Multiscale Spatio-temporal Dynamics of the Global Internet

Renesys: Andy Ogielski
Jim Cowie
AT&T Labs: Walter Willinger
Anja Feldmann (Univ. Saarbruecken)
Dartmouth: David Nicol
Princeton: Ingrid Daubechies

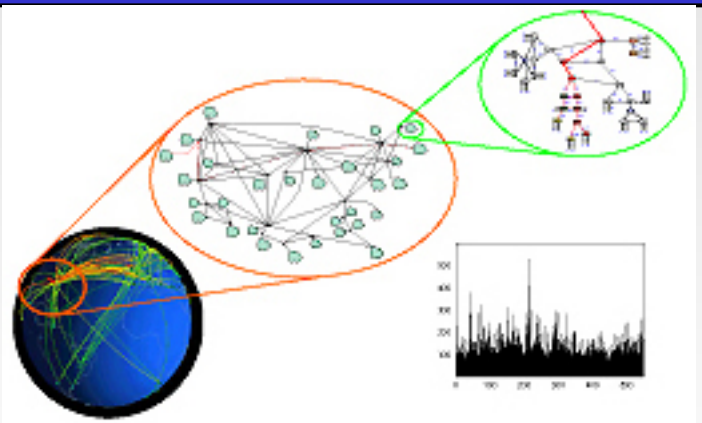

Goal: **Global Network Situational Awareness**

To watch stability, security, performance need multiple dynamic views on global networks, with explanatory and predictive powers.

Challenges: Need revolutionary advances in collecting, integrating, and representing extremely large, diverse, distributed network data streams.
Test with scalable, complex topology Internet modeling & simulations.

New ideas: Invent and develop novel multi-resolution analysis (MRA) tools, analyze multi-node real and simulated networks side-by-side:
- across multiple scales of time, topology and protocols
- focus on structural mechanisms of complex behaviors in Internet.

First broad investigation of the spatial and temporal patterns of global routing and traffic flows

Multiscale Spatio-temporal Dynamics of the Global Internet

Recent results

Powerful new multiresolution analysis techniques developed:

- Correlating traffic data sources measured at different granularities

- Multiscale aggregation of topology and traffic flows

Advances in Model Construction and Simulation:

- BGP routing convergence modeling & simulation studies

- Discrete event fluid modeling of individual TCP flows, fully interactive, validation, mixed packet/fluid capabilities

- SSFNet and SSF in DARPA, DoE and DoJ programs and in industry

Multiscale Spatio-temporal Dynamics of the Global Internet

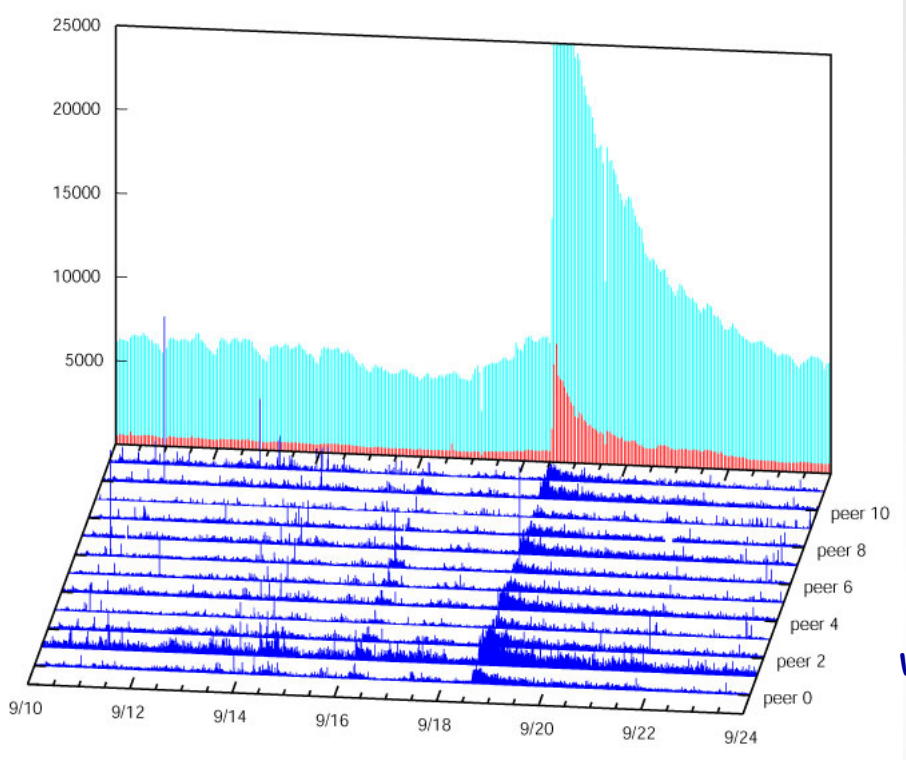

Latest news:

Renesys team discovers global routing instabilities caused by Code Red 2 and Nimda worm attacks .

Based on analysis of BGP message streams from 100+ border routers worldwide.

www.renesys.com/projects/bgp_instability

18 September 2001 example: big surge of route withdrawal messages sent by BGP routers from 12 major Internet providers (blue) . This routing instability correlates with attack rate by **Nimda worm** probes (red) and rate of MS host infections (cyan).

Note no routing instabilities on September 11 – this was a local effect.

Preliminary analysis: various edge router failures triggered by exponential worm spread cascade into global routing instabilities lasting **many hours**.

Focus of today's presentation

Closing the loop:

from global-scale Internet measurements

to

global-scale Internet modeling & simulations

Multiscale Spatio-temporal Dynamics of the Global Internet

Global Routing Instabilities

Microsoft worms and **BGP storms**

July 19: *Code Red 2 storm*

Sept 18 –19: *Nimda storm*

Credits

- Andy Ogielski and Jim Cowie, Renesys
- Brian Premore and Yougu Yuan, Dartmouth & Renesys
- Worm traffic data from several /16 networks courtesy of Vicki Irwin (SANS Institute), Ken Eichman (CAS), Vern Paxson (ACIRI).
- RIPE NCC Routing Information Service, Amsterdam
  data from 6 European collectors, 161 BGP sessions

Multiscale Spatio-temporal Dynamics of the Global Internet

AS (Autonomous System):

A collection of networks under the same administrative control

-Examples: universities, corporations, other institutions

-Current total active ASes: 11,000 (over 22,000 assigned)

BGP (Boundary Gateway Protocol) :

Glues together ASes, enabling communication between any two locations in the Internet

- BGP is executed by routers within ASes
- Each AS has from 1 to 1000+ routers running BGP
- Current total BGP routers: over 100,000 (estimate)
- Packets forwarded based on address *prefix*
- forwarding based on route announcements/withdrawals
  - a withdrawal means the router has no way of delivery to that prefix!

Global Internet Routing Instabilities

"Catastrophic instabilities are expected behavior of large engineered systems"

Qualitative features

rate	duration	diversity
High rates of route changes: • magnitude • acceleration • variance	Very long times: • long relative to baseline noise • long relative to expected routing table convergence time	Seen at many observation points: • many external BGP peers • many exchanges • Intra-AS networks Seen in high diversity of routing traffic content: • number of prefixes • number of routes

Case Studies

Code Red v2 attack

Nimda worm

prefix announcement rate in 30 sec intervals

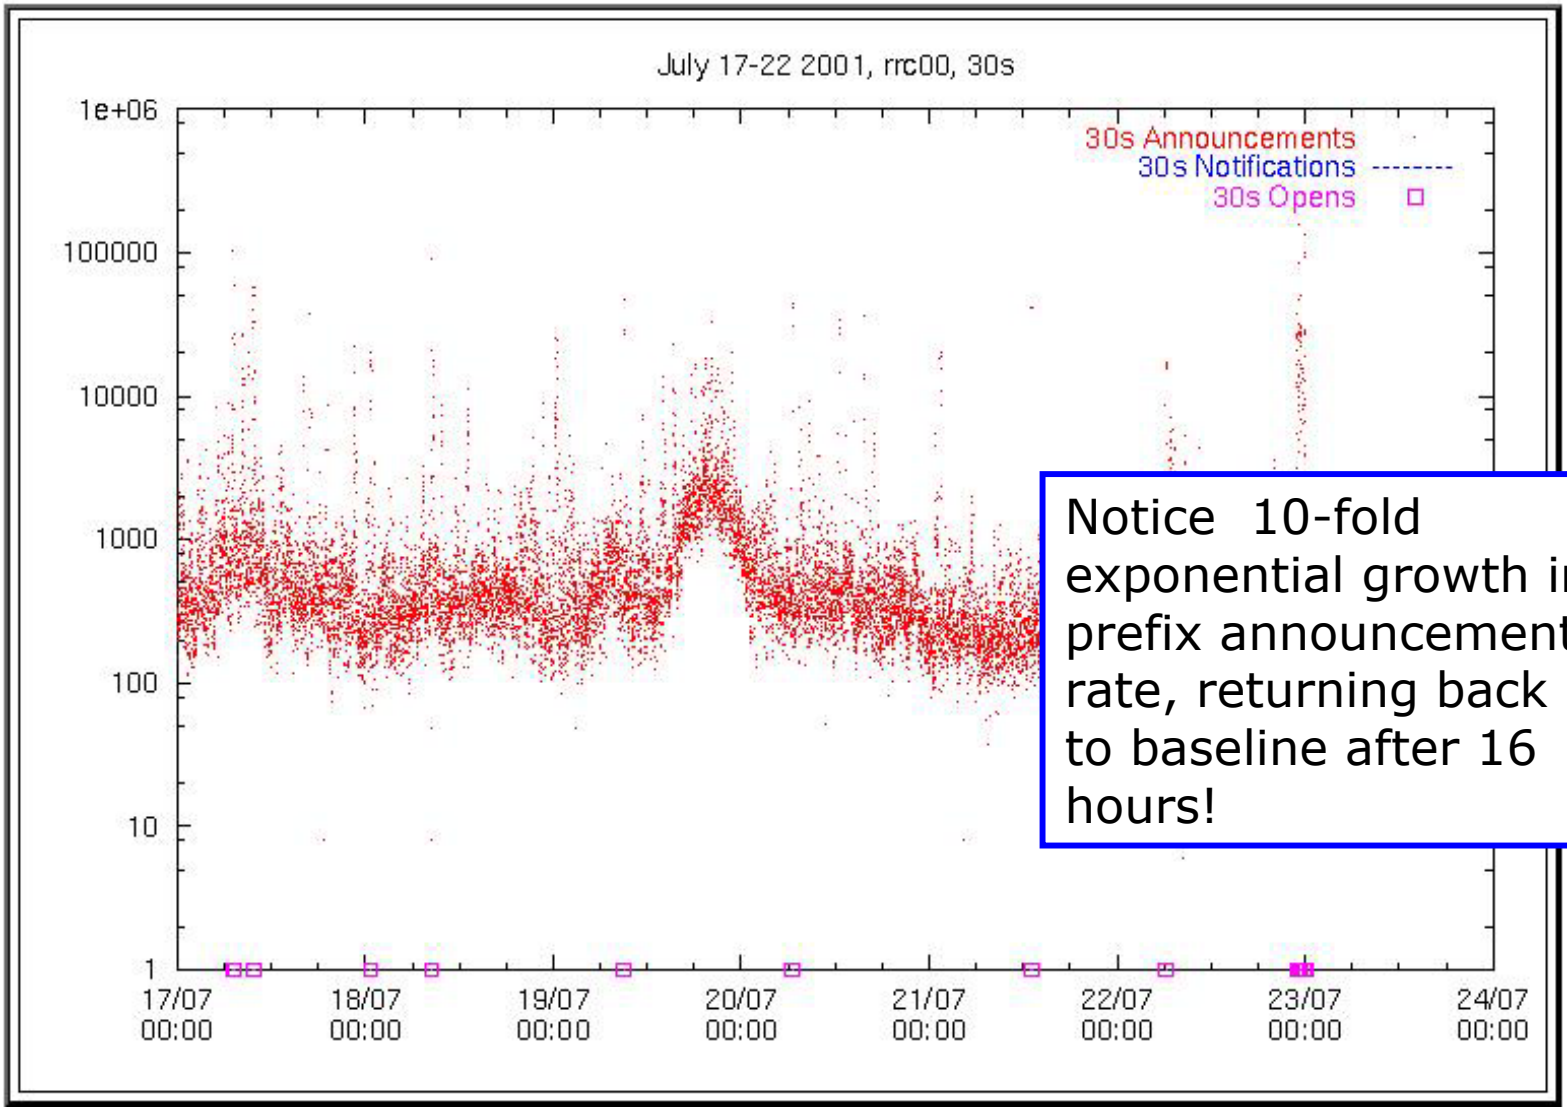

Notice 10-fold exponential growth in prefix announcement rate, returning back to baseline after 16 hours!

renesys

What is going on?

Look for behavior across peers

Look for behavior across origin ASes

Look for behavior across prefixes

Look for route withdrawals by prefix length

Look for route lifetimes

Prefix withdrawals by peer

RIPE NCC, July ’01, 60-min intervals

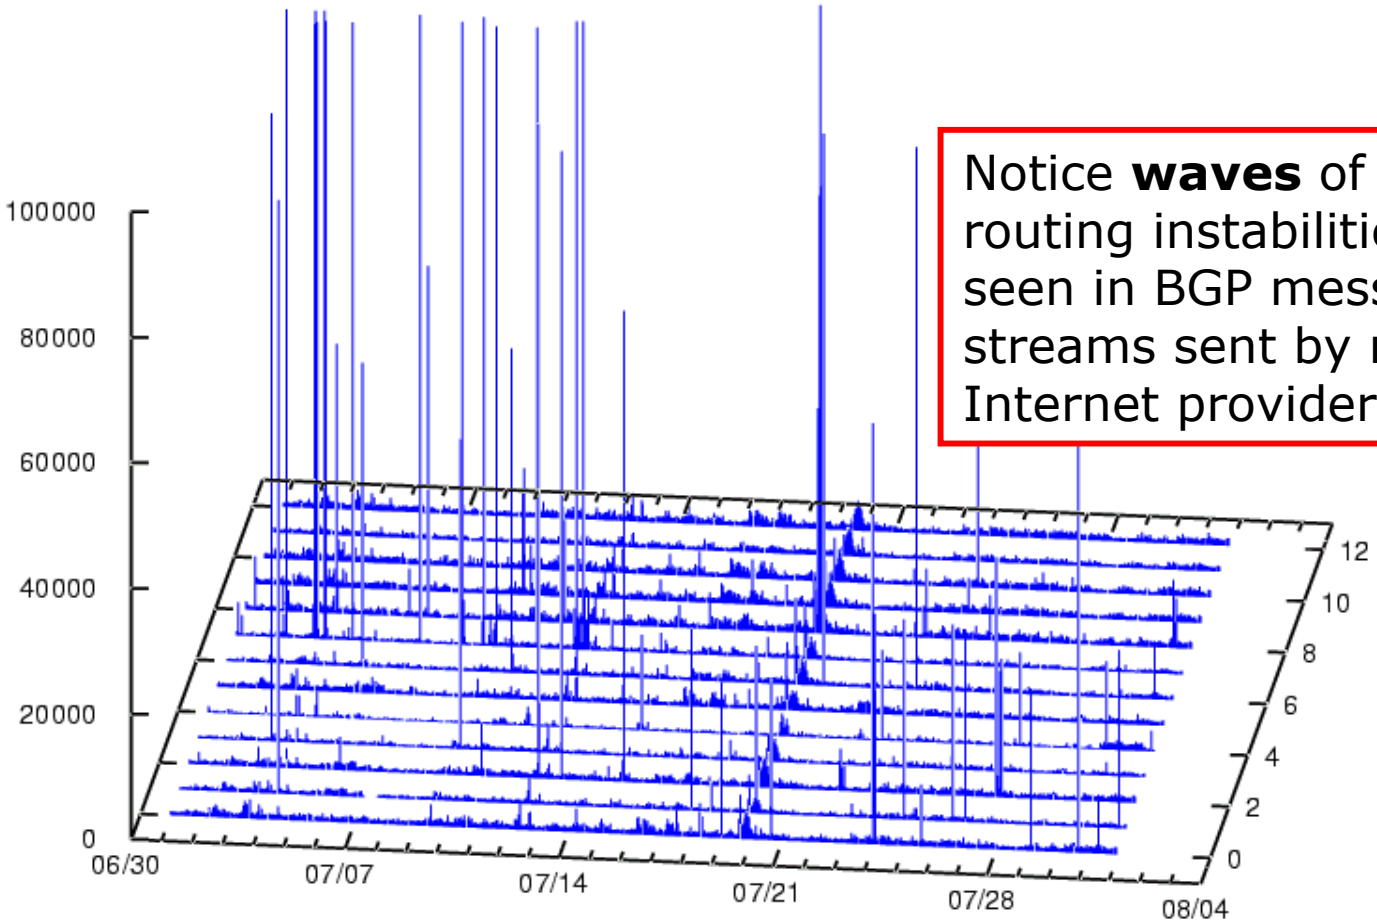

Notice **waves** of routing instabilities seen in BGP message streams sent by major Internet providers

renesys

Results of detailed analysis:

July 19 BGP long-term instability is <u>not</u> driven by a localized network failure:

- no suspect peers – all major ISPs show similar trend
- no suspect prefixes – most prefixes churn
- no suspect routes – most routes churn

Correlate with other Internet data

July 19 BGP storm correlates in time with Code Red II worm attack

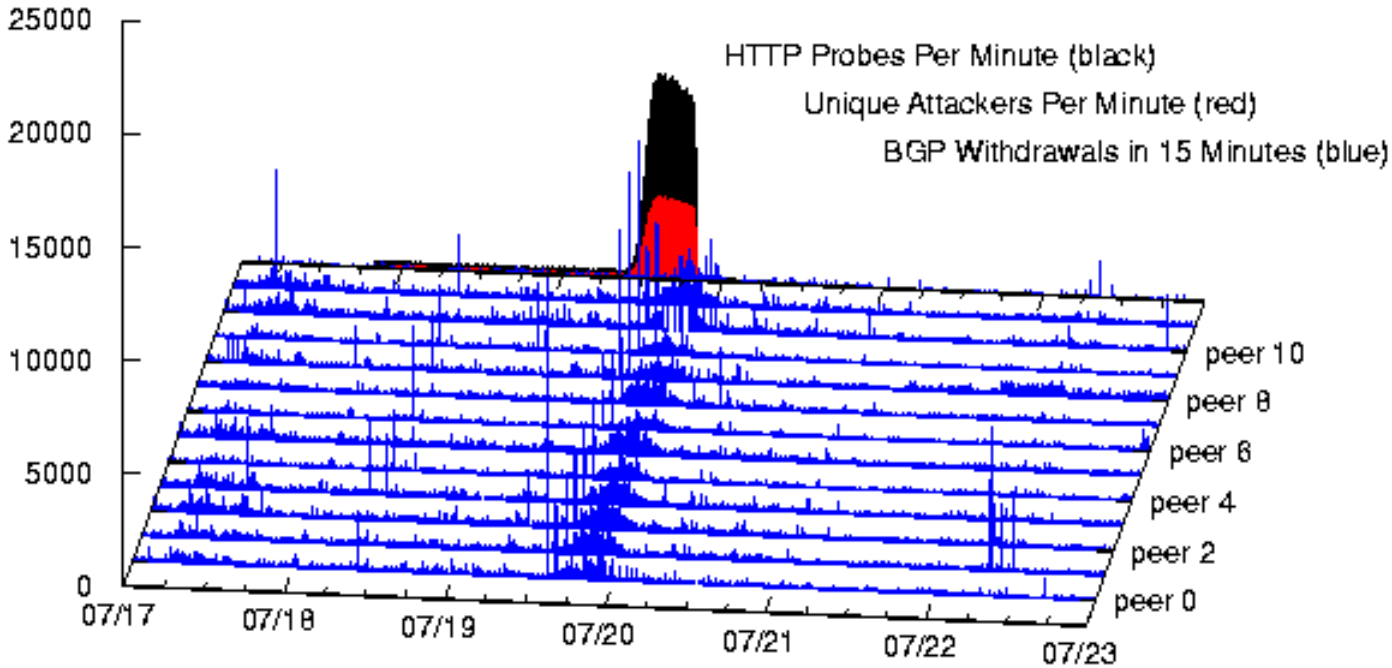

renesys

September 18 BGP storm correlates in time with Nimda worm attack

Smaller storms: leakage of malformed route announcements

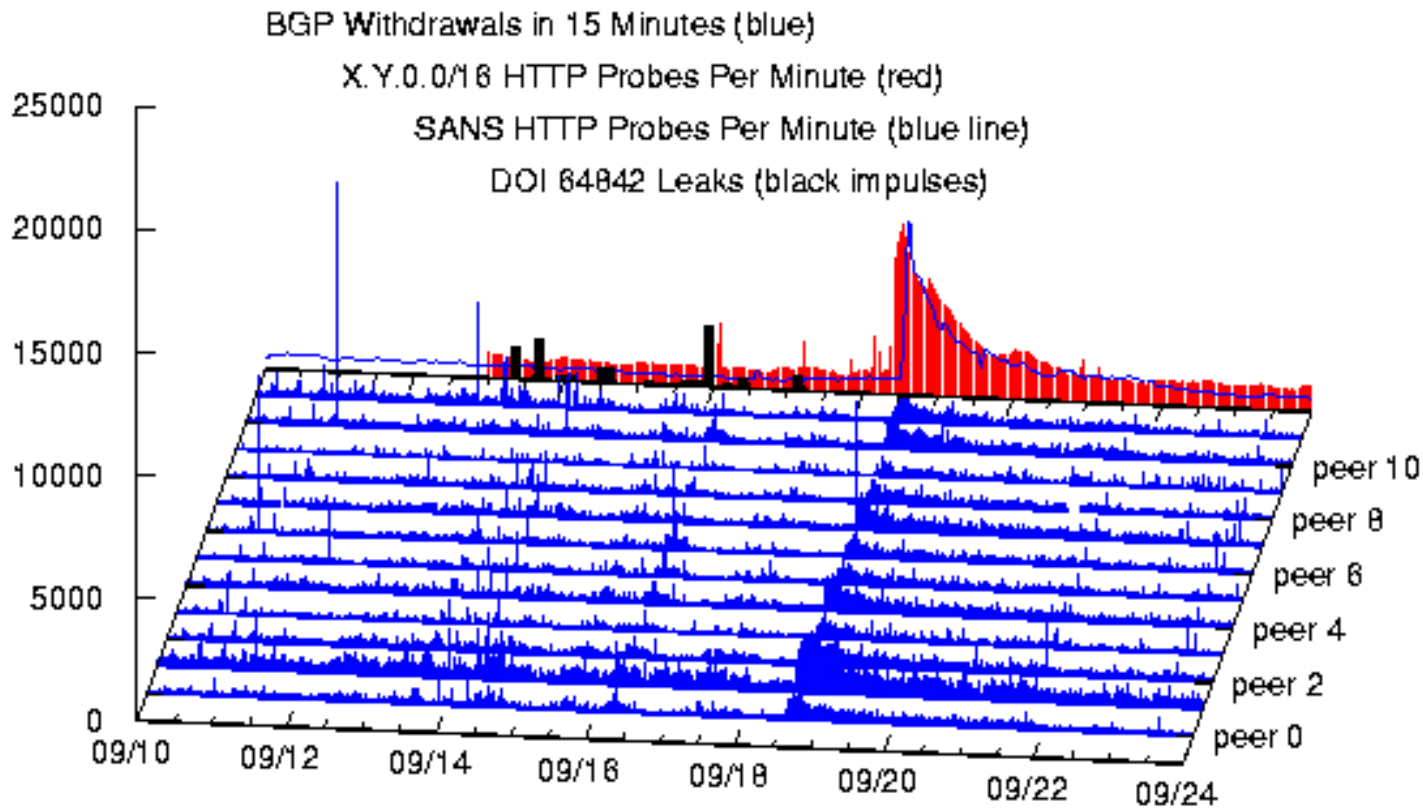

renesys

Worm-induced BGP storms seem to arise from BGP connectivity failures at very many locations.

BGP routers can fail due to:

- router CPU overload
- router out of memory, cache overflows
- router OS failures

Possible worm traffic causes:

- traffic intensity
- traffic diversity
- HTTP servers in routers (mngmt interfaces)
- failures in network gear (DSL routers,…)
- IGP (Intra-AS) flapping and routing failures
- proactive disconnection of networks

Closing the Loop : Simulate to Understand BGP Storms

We are uniquely positioned to investigate causal effect of worms on global routing infrastructure

- SSF and SSFNet provide mature simulation & modeling infrastructure
- We have developed a comprehensive BGP module and have on-going studies of BGP
- Model size is a key issue, PDES absolutely required, possibly also out-of-core techniques

Research Plan

- build a number of topology models based on Internet
- reconstruct conditions leading to cascading route withdrawals
- model router device behavior as a function of BGP traffic
- model worm propagation and traffic w.r.t. routing
- validate behavior, look for correlations in router failures and effect on traffic behavior

Understanding BGP Convergence Behavior

On-going study

- Brian Premore & David Nicol, Dartmouth
- Tim Griffin, AT&T Research

Goal is to understand the dependence of how quickly BGP "settles in" (for a while) on a route (for a given prefix) on various configuration parameters. Why? Lack of convergence degrades traffic delivery

Observation : BGP has a "heart-beat" w.r.t. advertised route changes

- frequency is a key parameter

Methodology of study

- Observe the phenomenon in the wild
- Model the system
- Run 10's/100's of thousands of experiments
  - project developed new software framework for experimental design and results management/exploration
- Look for sensitivity of convergence as a function of parameter Minimum Route Advertisement Interval (MRAI)

Stethoscope on the Internet

Histogram of route lifetimes
(period over which route does not change)
Y-axis is logarithmic

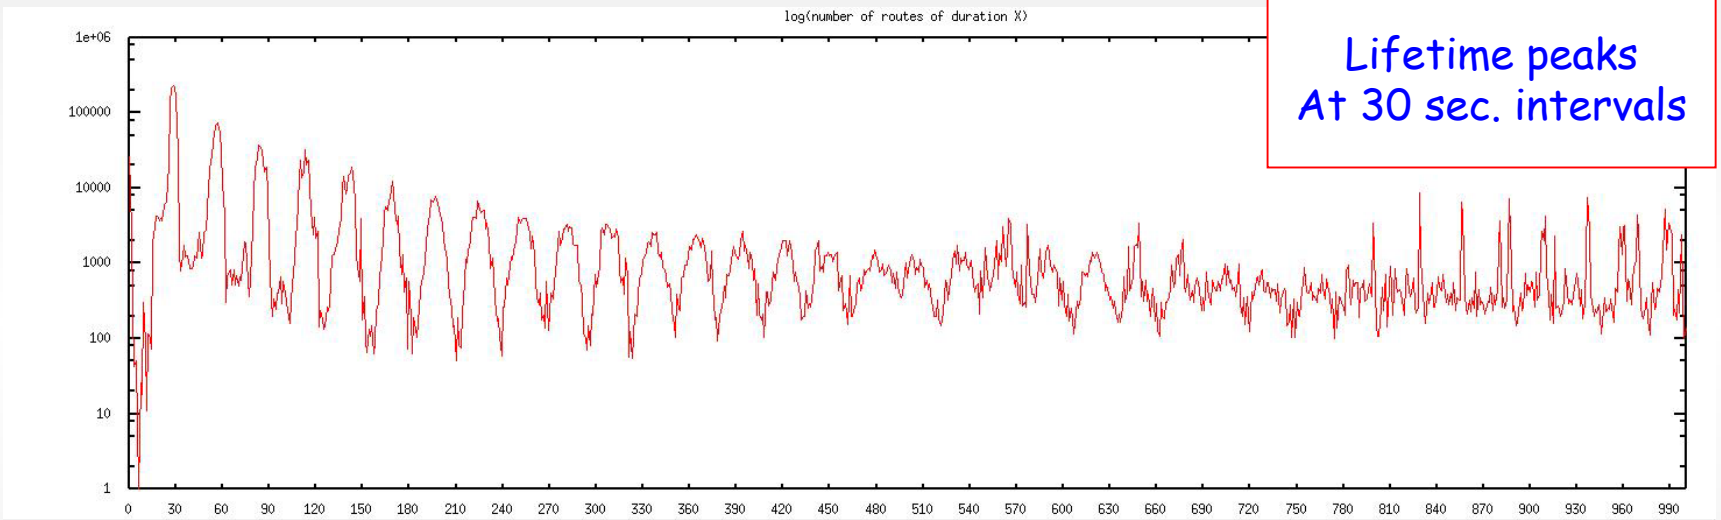

Lifetime peaks
At 30 sec. intervals

renesys

Experimental Design Space

A huge number of model parameters to explore:

-type of change injected
(announcement or withdrawal)
-network size
-network topology
-MRAI value
-router workload
-routing policy
-link delay

-protocol implementation choices
- rate limiting of withdrawals
- rate limiting granularity
- loop detection policy
- randomized tie-breaking
- timer jitter
- route flap dampening

over 10,000 different model configurations used
- scripts for organizing simulations *[Srdjan Petrovic]*
- used BGP module included with SSFNet

procedure:
- inject a change
- measure time required for convergence

Results of Study

Internet is running suboptimally : observed optimal values much lower than values used in practice!

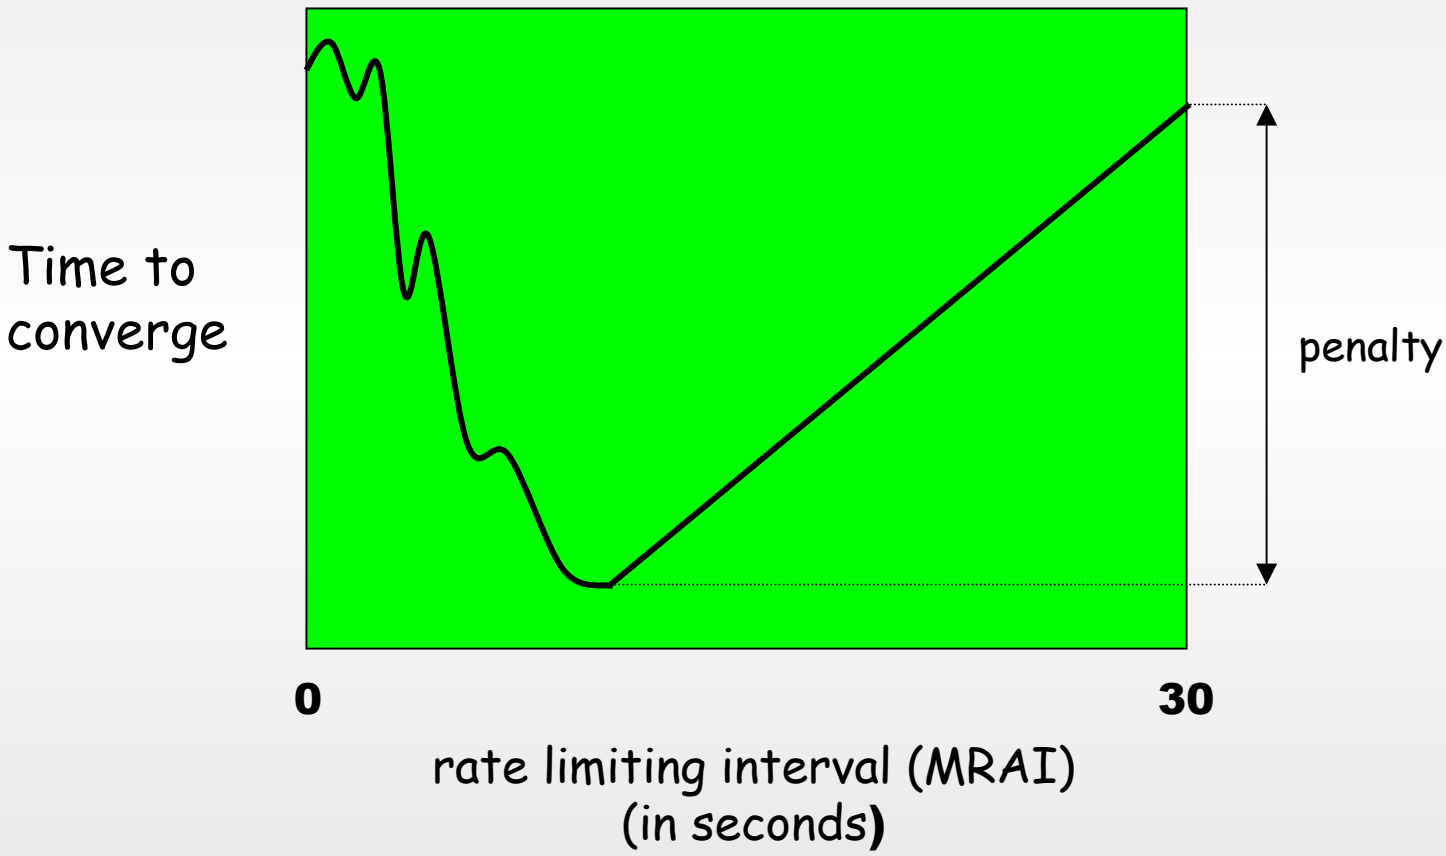

Summary

Our team is exporing Internet dynamics through a variety of means

Measurement of Global Internet

Methodology for seeking global effects from noisy detailed data

Modeling and Analysis of Routing Behavior

Tools for

- managing large-scale simulation studies
- simulating very large scale networks, many protocols

A spectrum of new capabilities enable us to investigate the vulnerability of the routing infrastructure